

SACR AI SOC MARKET LANDSCAPE 2025

FRANCIS ODUM
RAFAŁ KITAB

IN COLLABORATION WITH

 Qevlar AI  D3  CMU
ZERO  CROGL  torq
 prophet  sekoia  Mate  Dropzone AI
RADIANT  Exaforce  INTEZER  LEGION

About Software Analyst Cybersecurity Research

SACR is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000** readers and followers, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

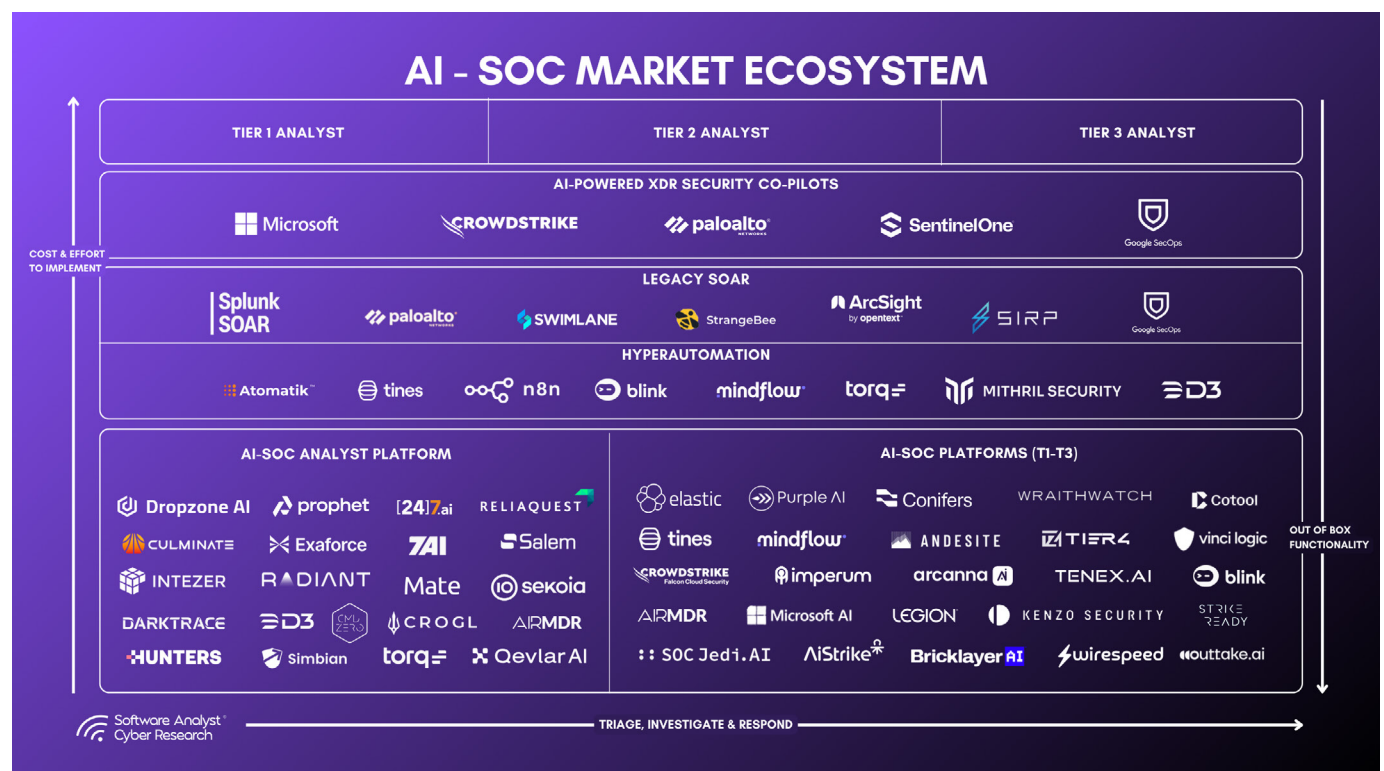
Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, or timely, opinionated insights help modern security leaders make smarter, faster decisions.

Table of Contents

Executive Summary	4
Introduction	5
The Modern SOC	7
Challenges With The SOC In 2025	8
AI Within The SOC	9
AI SOC Architectural Options	10
Operational Guide to AI SOC Adoption	14
Risks and Considerations with Agentic AI in the SOC.....	16
Evaluation Methodology	17
Radiant Security	21
Conclusion.....	23





Executive Summary

AI is moving from experimentation to execution in the Security Operations Center (SOC). Around the world, security teams are beginning to embed AI agents directly into detection, investigation, and response workflows, transforming how the SOC operates at scale.

This report captures SACR's latest perspective on this shift, based on extensive research, direct CISO feedback, and hands-on evaluations of leading solutions. It defines three core architecture models for AI-enabled SOC operations and maps 13 vendors to real-world use cases matched to varying SOC maturity levels.

This research report was grounded in seven technical benchmarks and informed by practitioner insight, the report delivers a practical decision framework and a phased rollout strategy to help leaders balance trust, automation, and operational impact. The goal is to equip security leaders with the clarity to identify the right entry points, reduce deployment risk, and accelerate measurable outcomes from AI in the SOC.

Authors

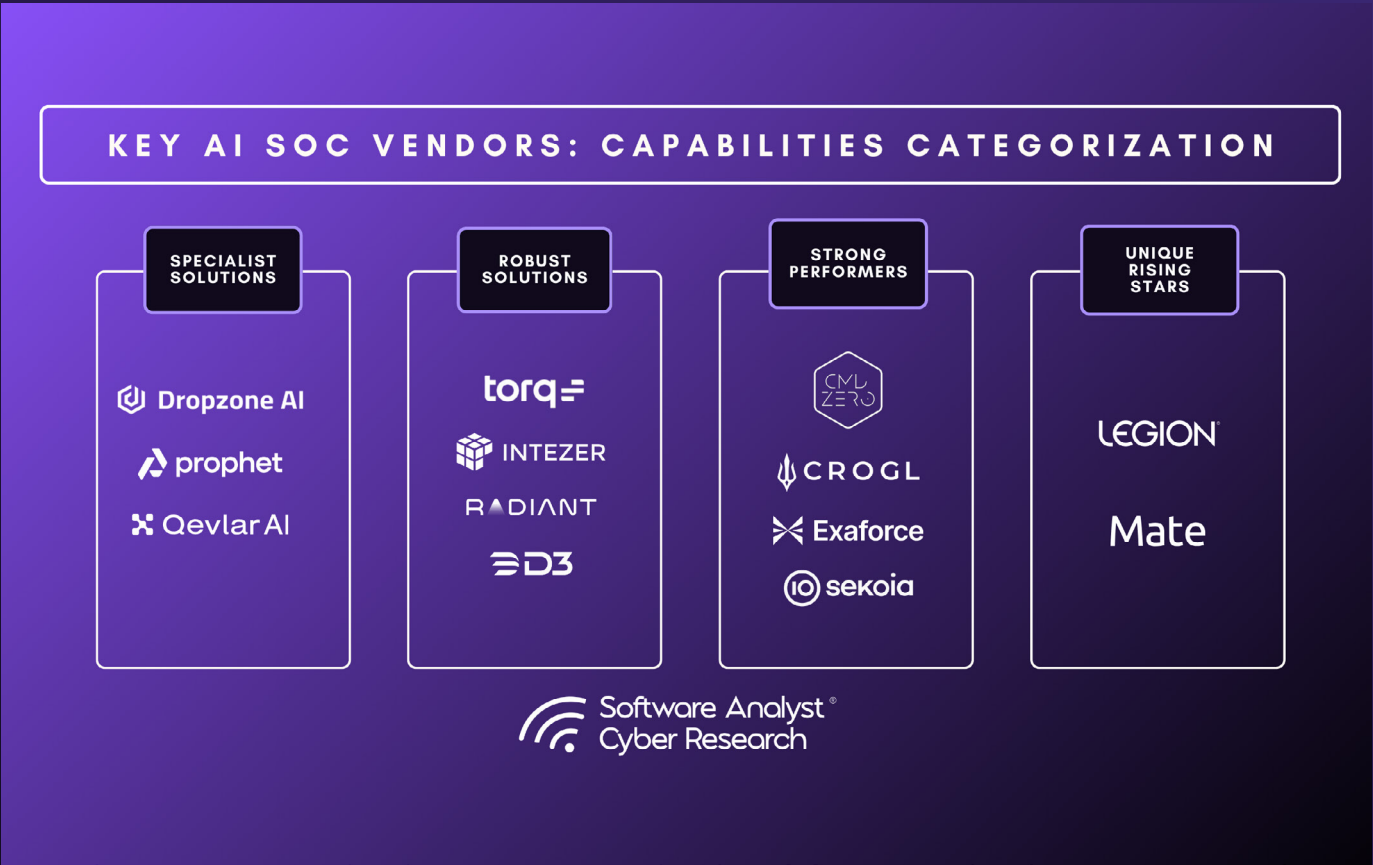
Francis Odum is the Founder/CEO of the Software Analyst Cyber Research, where he leads the firm's research and engagement with cybersecurity leaders.

Rafal Kitab is a SOC and Incident Response leader at ConnectWise with extensive experience working as a Security Analyst, Engineer, Architect, Incident Responder and recently a Director. He brings considerable experience in Security Operations and shares his first hand experiences.

Introduction

Security operations centers (SOCs) stand at a turning point in the rapidly evolving digital landscape of 2025. Among the CISOs in our network, this emerged as one of the year’s top strategic priorities. We conducted considerable research on this category last year — [Revolutionizing Security Operations: The Path Toward AI-Augmented SOCs in 2024](#). However, we felt the need to conduct a more comprehensive report this year to update our research on the latest findings as of 2025 since many things have changed.

One of the most significant changes is how fast companies have emerged. The ecosystem for AI SOC vendors has evolved over the past 12-15 months, with many different vendors pivoting to solve the problem. We’ve seen a surge of vendors entering or pivoting into the AI SOC space, each promising to alleviate alert fatigue, automate investigation, and augment analyst workflows. As the ecosystem rapidly expands, it’s become increasingly difficult for security leaders to distinguish between marketing claims and meaningful capabilities. The signal is buried in noise, and clarity is urgently needed. This report aims to provide the depth of clarity required.



The Challenges Remain

Security teams today are handling an overwhelming volume of alerts while also struggling with persistent staffing shortages. Enterprises now face thousands of alerts per day, often with limited capacity to investigate them in a timely way. Against this backdrop, AI has moved from an experimental idea to a working system inside the SOC. AI SOC platforms are now being used to monitor environments, investigate alerts, and respond to threats with less manual effort. These tools are helping reduce detection and response times, ease analyst burden, and improve coverage across increasingly complex environments.

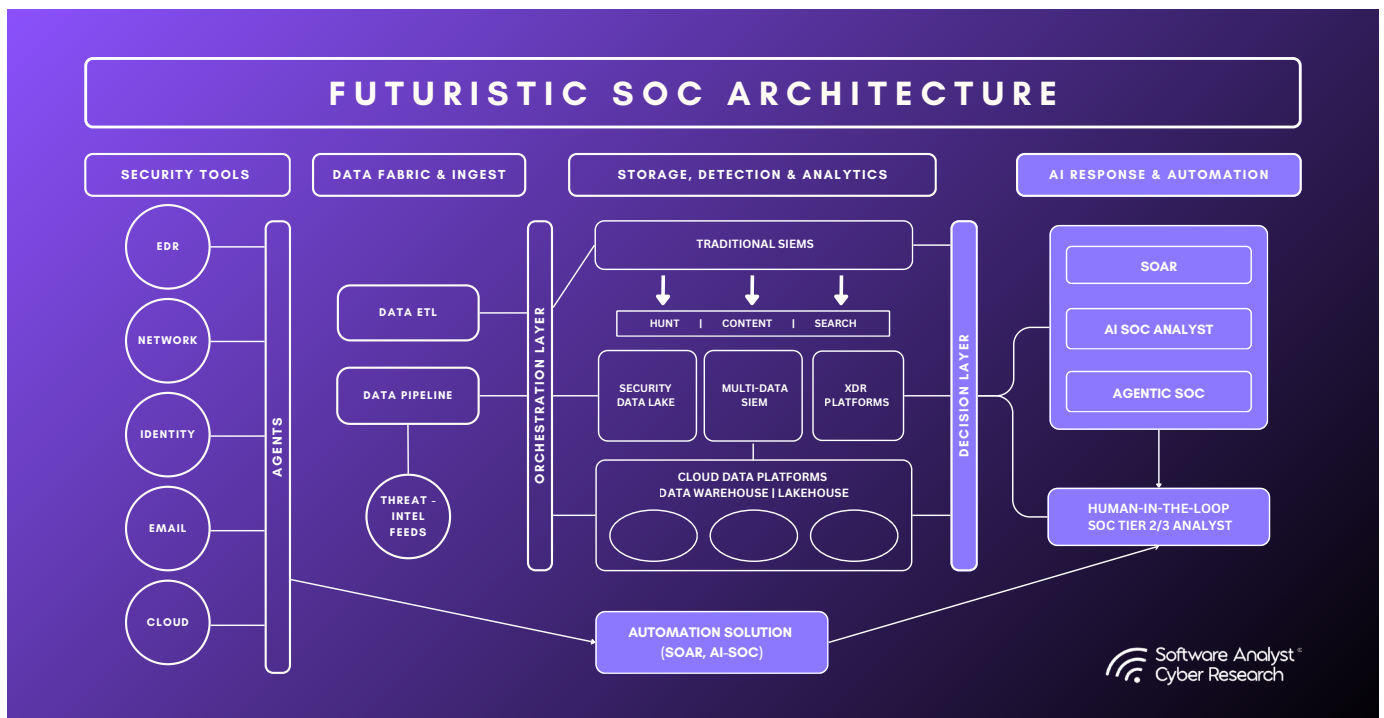
This shift raised a broader question: Is this the moment when security teams gain a lasting upper hand, or is AI simply another fleeting trend in the SOC? To answer that, we conducted the most in-depth report to date on this emerging category. The insights, analysis, and frameworks presented in this report aim to equip CISOs, security analysts, and organizational decision-makers with the tools to discern hype from substance and make informed strategic choices for the future of their cybersecurity posture.

To ground our research, we partnered with 13 prominent AI SOC vendors, examining their approaches to automation, alert triage, investigation, and response through rigorous benchmarking and real-world evaluations. We sent each vendor a detailed questionnaire and asked them to walk us through how their platform works. We then used their answers to compare and map capabilities using SACR's internal assessment framework. The result is a structured view of how these platforms differ in architecture, automation models, and overall strengths.

There are many other competitors on the market we have not yet evaluated but plan to include in future research. Some of the vendors not featured in this edition but actively building agentic solutions for the SOC, include, Tines, SevenAI, Elastic AI, AStrike, Bricklayer, Simbain, and others.

The Modern SOC

As security environments grow more complex, the modern Security Operations Center (SOC) has evolved into a layered architecture, where each tier plays a critical role in managing scale, ensuring context, and accelerating response. We break it down into three layers (with a spotlight on the right side today):



Data Fabric Layer

In SOC teams, a layer of raw security data from diverse sources must be standardized and processed for effective analysis before being sent to a SIEM (or automation solution). This category has two distinct types of vendors. 1) Those focused on building and managing the engineering data pipeline. 2) Those prioritizing data filtering and enrichment to improve detection quality. The key task is formatting data into a consistent structure to enable seamless integration and adding contextual information, such as IP geolocation or threat intelligence feeds, to improve data quality. We'll uncover how this layer increasingly intersects with automation efforts.

Storage and Detection Layer

Companies leverage threat detection rules to define the logic for identifying malicious activity in their data. Once processed, data is routed to a centralized repository for storage and analysis. This may involve a SIEM (Security Information and Event Management) offering real-time monitoring and alerting or a cloud-based data lake designed to reduce costs.

AI Response & Automation

When detection rules are triggered, or alerts are generated, SOC analysts thoroughly investigate these alerts, assess their severity levels, and implement appropriate remediation measures. Modern SOC automation solutions are evolving to adopt a more proactive approach, integrating directly with security tools rather than relying solely on SIEM alerts. This advancement allows for enhanced alert enrichment and contextual analysis, leading to more efficient remediation processes. Analysts can now differentiate real threats from false positives more quickly and conduct thorough assessments and containment strategies. For incident response, teams can now manage security incidents more effectively, conducting deeper investigations by leveraging indicators discovered during the triage phase, with legacy SOAR solutions proving particularly effective in this domain. While AI technology is poised to revolutionize SOC operations in this layer, its adoption remains complicated by SOC leaders' ongoing concerns about SIEM-related costs and implementation challenges.

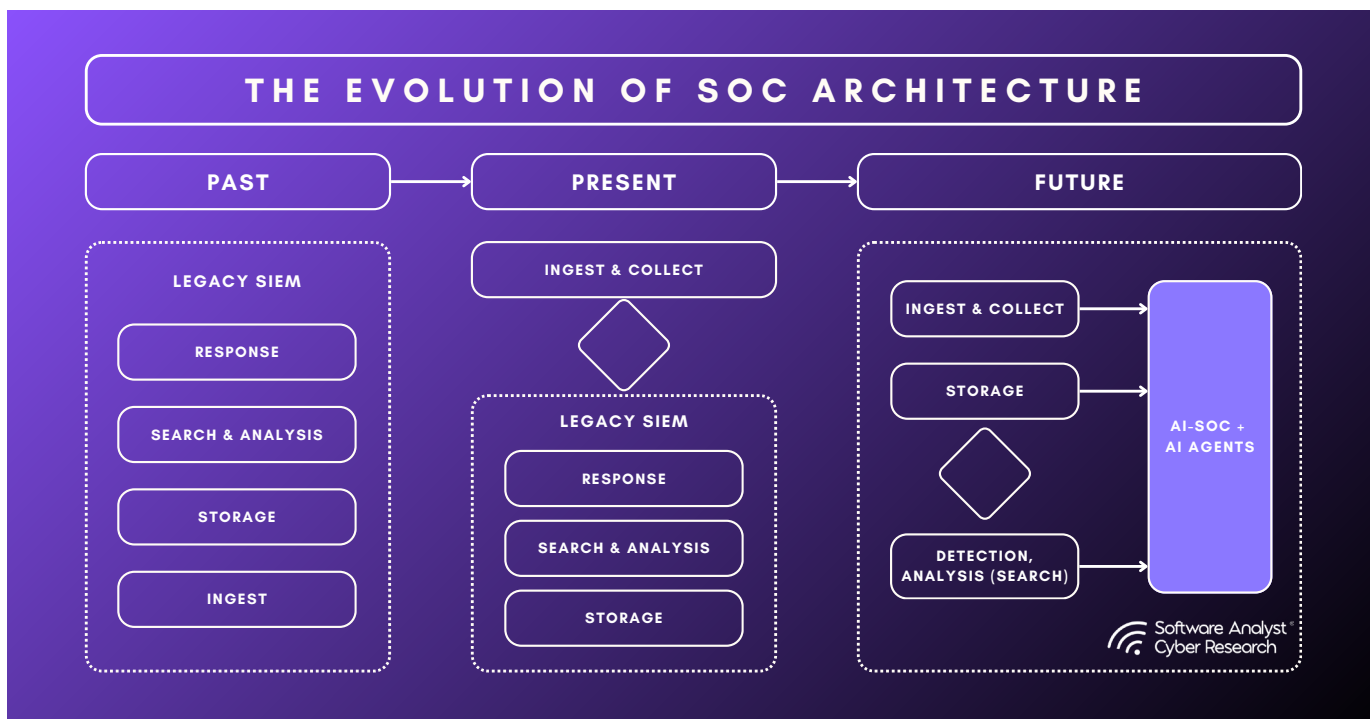
Challenges With The SOC In 2025

According to a survey of 300+ CISOs:

- **The scale of the alert tsunami is untenable:** Across the 282 organizations surveyed, teams face ~960 security alerts every day, and enterprises with 20k+ employees are drowning in more than 3k alerts daily, generated by an average of 28 different tools. Analysts openly describe “way too many data sources” and a “tsunami of data” as issues that plague every layer of the SOC.
- **Alert fatigue has become a systemic risk:** On average, 40% of all alerts are simply never investigated, and 61% of security teams admit they have ignored alerts that later proved critical, exposing customer data, taking systems offline, or driving direct business losses. Compounding the danger, the mean time to investigate sits at 70 minutes while phishing-based breaches can succeed in <1 hour, giving adversaries a decisive head-start.
- **Resource constraints are forcing teams to accept blind spots:** 57% of organizations now suppress detection rules just to keep workloads manageable, and the first rules to be disabled are in cloud and identity, the two fastest-growing attack surfaces. Leaders acknowledge this “necessary risk,” underscoring how urgently they need automation that can restore visibility without ballooning head-count.
- **Market sentiment has already shifted toward AI:** “AI for Security” has vaulted into the top-three priorities for CISOs, behind only data and cloud security and 88% of organizations that do not yet run an AI-driven SOC plan to evaluate or are actively starting one up within the next 12 months. The market is moving from if to how fast AI can be integrated into everyday SOC workflows.
- **AI is expected to own the majority of SOC workload within three years:** Security leaders project that AI platforms will shoulder ~60% of all SOC tasks by 2028, with 83% believing AI will handle at least half of operations. The metrics they will watch, MTTI, MTTR, and 24/7 coverage, align directly to AI’s strengths, signalling a near-term, ROI-driven buying cycle for purpose-built AI-SOC solutions.

AI Within The SOC

Last year, we wrote about the evolution of AI in the SOC. At the time, it was more an exploratory report into the category. In recent months, we've validated the problem through conversations with major security leaders globally. As SOC's grapple with escalating alert volumes, rising costs, and mounting pressure to reduce detection and response times, artificial intelligence has emerged as a necessary augmentation. It is not a replacement for human analysts, but a tool to help them manage growing operational complexity.



This second version of our report shifts the focus from exploring the need for AI in SOC operations to evaluating which AI-enabled vendor solutions are truly capable of enhancing analyst productivity, improving visibility, and reducing response times.

Recognizing ongoing skepticism around the transparency, efficacy, and adaptability of AI tools, this report segments the vendor landscape into four defined categories:

- AI-powered XDR co-pilots
- Automated Tier-1 alert analysts
- Advanced AI-driven threat hunting platforms
- Workflow-centric AI automation engineers

We apply rigorous benchmarks including measurable reductions in mean time to detect and respond (MTTD and MTTR), improvements in

false-positive management, and operational cost-effectiveness to guide SOC leaders in selecting partners that align with their specific environments. This report also puts a spotlight on the use of AI agents within Security Operations Centers. SOC's typically rely on systems that take data as input and generate alerts based on predefined logic. As the volume of incoming data grows, so does the risk of alert fatigue. Agentic AI is designed to reduce that burden. By autonomously handling large volumes of alerts, it helps analysts focus on the events that truly require their judgment and expertise.

We examine how this technology is being applied across two types of environments:

1. Internal SOC's operated in-house
2. External SOC's managed by service providers

AI SOC Architectural Options

We reviewed many architectural breakdowns on the market, including those written by Andrew Green.

Functional Domain: What does it automate?

Automation/Orchestration (SOAR+) & Agentic SOC

These platforms are designed to be the “central nervous system” of security operations, coordinating and automating responses across a wide range of security tools and data sources. They go beyond simple playbooks, leveraging agentic AI to intelligently sequence actions, enrich alerts, trigger containment or remediation, and handle case management, often without human intervention. Their greatest strength lies in their ability to orchestrate complex, cross-tool workflows (SIEM, EDR, cloud, ticketing, etc.) at scale, using both rules and dynamic agent logic. This results in dramatic improvements in response speed, efficiency, and consistency, particularly for large, complex environments or MSSPs. These platforms can be user-defined (built from modular blocks) or enhanced with pre-packaged agentic functions, offering the flexibility to evolve as new threats and tools emerge.

Pure-play Agentic Alert Triage Platform

While orchestration platforms focus on end-to-end workflows, this category tackles a more specific challenge: reducing the noise and burden of alert overload. These platforms rapidly triage, classify, and escalate only the most relevant threats, serving as the “first line of AI-powered defense.” They ingest high volumes of alerts from existing security systems and autonomously filter out false positives and routine events. What differentiates them is their emphasis on agentic reasoning, applying learned behaviors, contextual data, and even large language models to determine alert priority and next steps. The value is immediate: analysts are freed from “alert fatigue,” and only cases truly requiring human expertise are escalated. These are ideal for organizations looking to boost SOC productivity without a wholesale change to their architecture.

Analyst Co-Pilot/Investigation

Platforms in this domain act as “digital teammates” for human analysts, offering on-demand assistance for investigation and decision-making. Rather than automating entire workflows, these tools focus on augmenting analysts’ capabilities in real time. This can include natural language chatbots that answer questions, generate queries, summarize evidence, or suggest next steps; or more advanced reasoning engines that assemble context and walk analysts through complex incidents. What sets these apart is their role as a bridge between human expertise and machine efficiency; they’re not here to replace analysts, but to make them faster, more accurate, and less prone to error. These platforms are particularly valuable for Tier 2/3 analysts and for organizations that want to scale knowledge without losing human judgment.

Workflow/Knowledge Replication

This is the cutting edge of “institutional memory” in the SOC. Workflow/knowledge replication platforms observe, record, and learn from how the best analysts operate, then turn those behaviors into scalable, repeatable automation. Often browser-based or using workflow capture technology, these tools create digital “twins” of expert processes that can be replayed across future incidents, training new analysts and driving consistency. The unique differentiator is their ability to codify not just what to do, but how and why it’s done, preserving tacit knowledge that would otherwise be lost to turnover or scaling. They’re a powerful answer for organizations with a few superstar analysts, or for anyone seeking to operationalize best practices across distributed teams.

Implementation Model (How is it delivered?):

User-Defined/Configurable (Deterministic, agent-building, low-code)

These solutions put the power in the user’s hands: they are toolkits or platforms that let SOC teams design, customize, and continuously tune the automations, workflows, or agents that drive their security operations. Using visual interfaces, scripting, or low-code builders, users can define detection rules, orchestrate workflows, build custom agents, and adapt the platform to unique organizational requirements. This approach maximizes flexibility, adaptability, and ownership, making it ideal for organizations with mature teams or complex, evolving environments. The trade-off is that these solutions require a higher level of expertise and ongoing maintenance but the payoff is a SOC that truly fits the business’s needs.

Pre-Packaged/Black-Box (No/limited customization, R&D-driven agents)

In contrast, pre-packaged or “black-box” solutions are delivered as ready-to-run platforms with minimal end-user customization required. The underlying logic, agents, or workflows are designed and maintained by the vendor, often drawing on extensive R&D, threat intelligence, and industry best practices. This model is all about rapid time-to-value: organizations can deploy advanced AI-driven SOC capabilities quickly and easily, without the need for internal development. The trade-off is reduced flexibility; users are largely limited to the capabilities and workflows provided “out of the box.” These solutions are perfect for teams that want to modernize fast, value ease of use, or lack the bandwidth for complex customization.

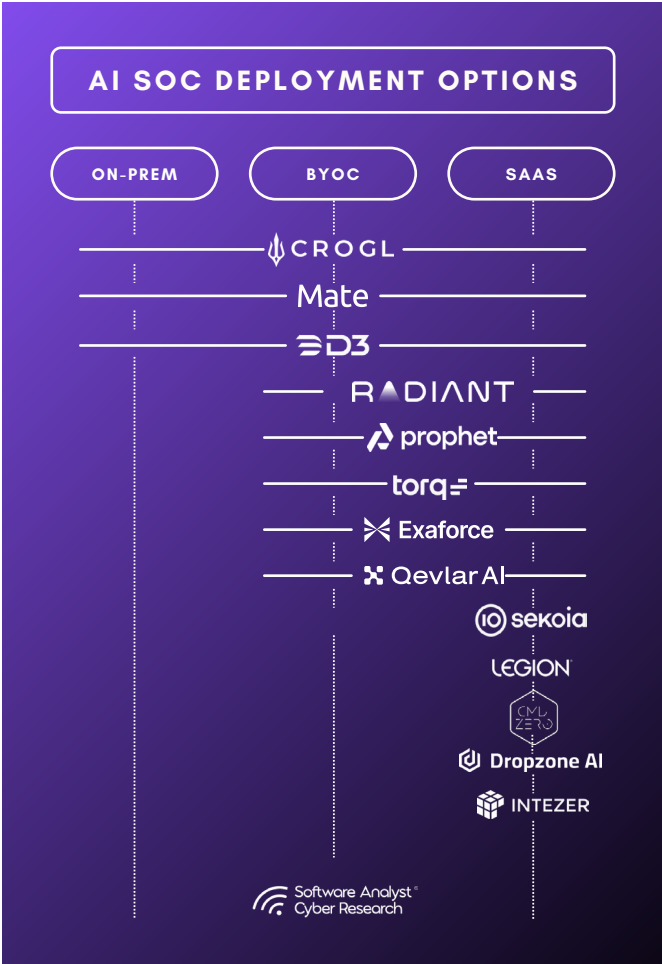
Deployment Options Amongst Vendors

Beyond architecture and configuration, deployment model is another key axis of differentiation. AI SOC platforms vary in how and where they can be deployed, shaped by performance needs, regulatory requirements, and cloud readiness. The

most common deployment model observed in this research was Software as a Service (SaaS), where the platform is hosted by the vendor and accessed over the internet. Some vendors also offer a “Bring Your Own Cloud” (BYOC) option, which lets clients use their own cloud infrastructure to store data and run the AI SOC platform on top of it.

Another notable deployment model, somewhat unique to AI SOC platforms, is support for air-gapped, on-premises environments. This option is particularly valuable for organizations with strict security or regulatory requirements, as it allows them to run the platform in complete isolation from external networks.

While all vendors demonstrate significant advancements in threat detection and response, the methods by which these improvements are achieved vary considerably. One of the most notable distinctions among solutions lies in their



underlying architecture. Through our evaluation of multiple offerings, we identified three primary architectural approaches, each with its own advantages, limitations, and implications for how AI is deployed within the SOC.

Connected & Overlay Model On An Existing SOC (SIEM)

This model refers to AI SOC solutions that are deployed as a layer “on top of” an organization’s existing security stack. These platforms are delivered as cloud or SaaS services, and their integration into the customer environment is achieved primarily through APIs.

They do not attempt to become the central data repository or replace core SIEM/logging infrastructure. Instead, they ingest alerts and telemetry from tools like SIEM, EDR, cloud, and identity sources, then apply automated enrichment, reasoning, or response logic before handing results back to the SOC team or case management system.

Their main appeal is rapid time-to-value. Because they do not require full-scale data migration, heavy tuning, or infrastructure build-out, they can often be deployed in days or weeks. These platforms are ideal for organizations looking to enhance investigation quality, automate triage, or add a layer of AI decisioning without disrupting their existing security architecture.

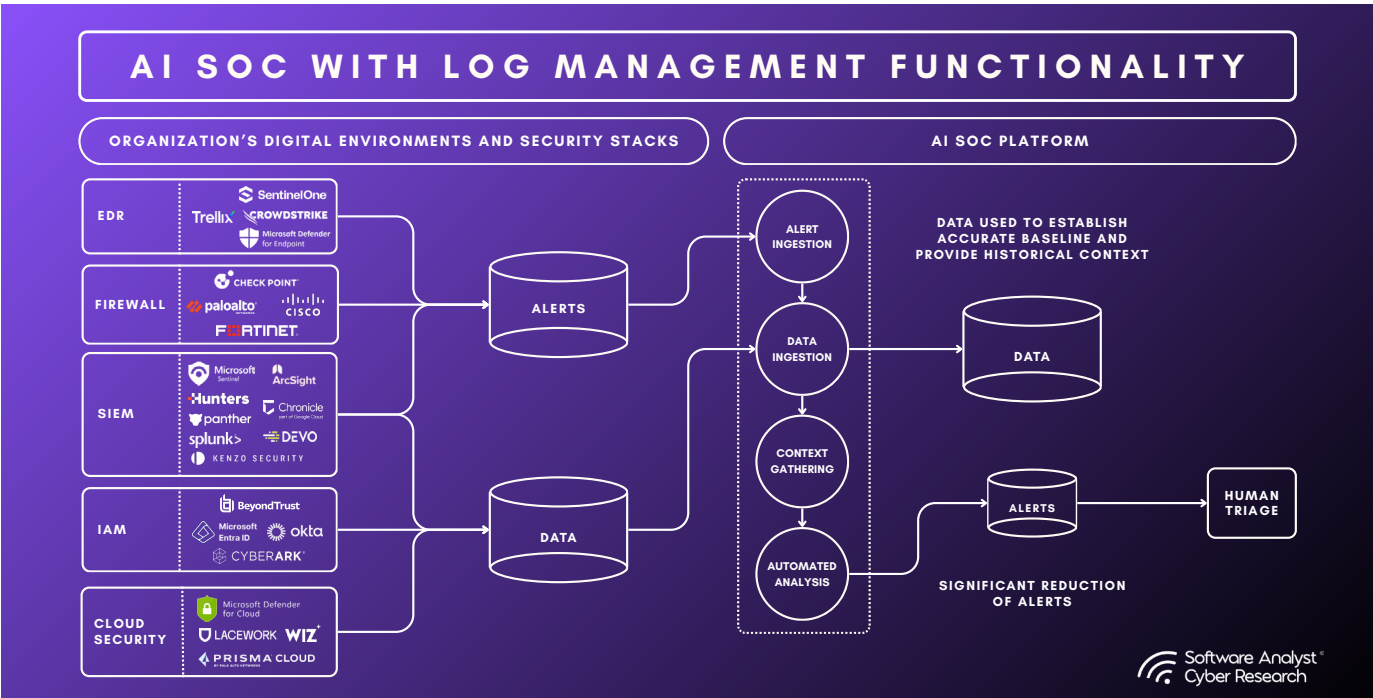
The trade-off is that these solutions rely on the fidelity of alerts and data generated elsewhere, they are only as good as the signal they are fed. They also tend to have limited behavioral analytics or anomaly detection capabilities, since they rarely have access to the full raw data stream.

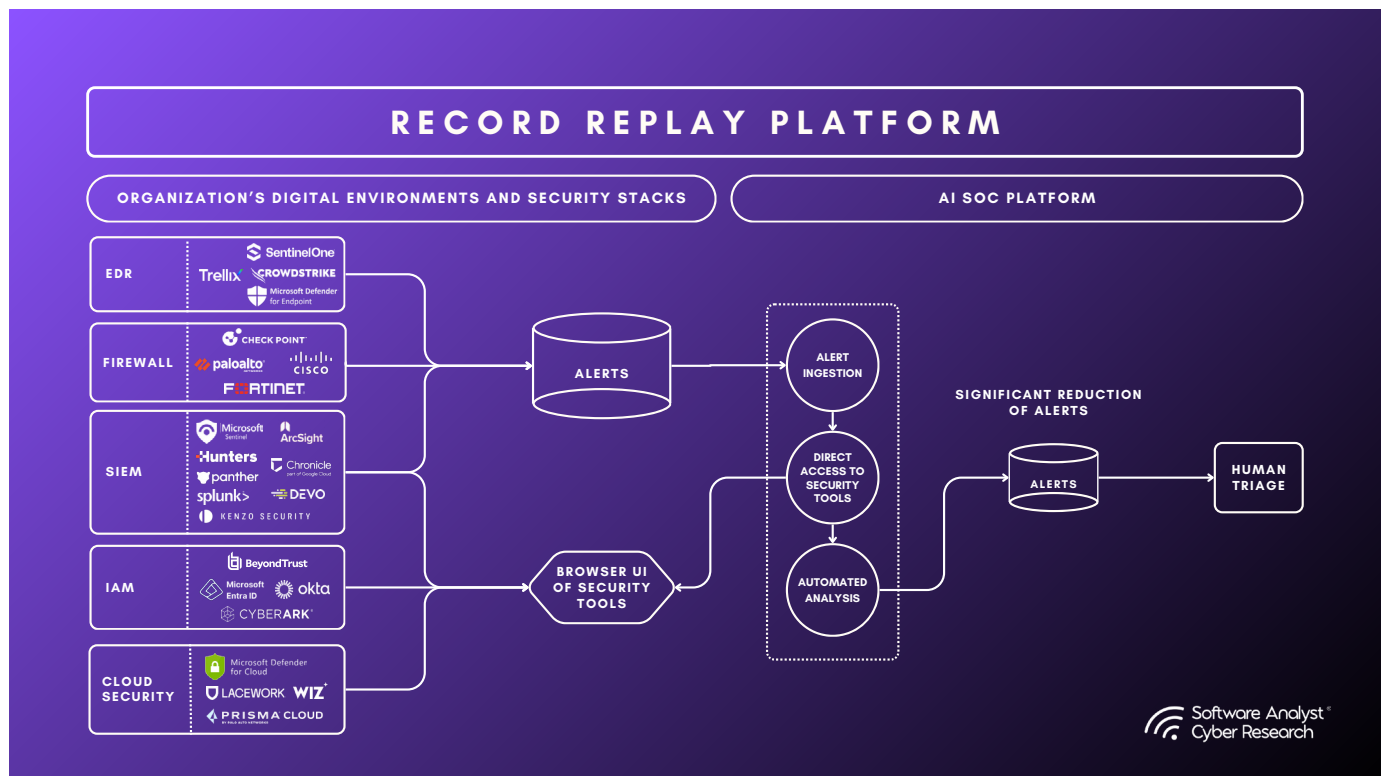
Integrated AI SOC Platforms

These platforms take a deeper approach to integration by ingesting, storing, and analyzing security data directly. In many cases, they act as a lightweight SIEM or even a full SIEM alternative depending on the use case. Unlike overlays, they access and retain raw logs and telemetry over time, which allows for more advanced behavioral analytics and long-term anomaly detection.

The key advantage is greater visibility and analytical power. By storing data internally, these platforms can establish historical baselines, surface subtle trends, and support retrospective investigation that is not possible with overlay-only models. Many also offer cost-effective log storage and retention, which helps reduce the high costs typically associated with traditional SIEMs.

These platforms are often hybrid in nature. They can act as a log storage offload or enrichment layer for organizations with expensive or overloaded SIEMs, while also serving as standalone detection and response hubs for smaller teams.





The trade-offs include higher operational complexity, the potential for vendor lock-in since the data resides within the vendor's environment, and additional security or compliance considerations, particularly for organizations with strict data residency or privacy requirements.

Human & Browser-based Workflow Emulation Platforms

This model represents the most human-centric and experiential approach. Rather than ingesting alerts through APIs and logs, these platforms capture, learn, and replicate the investigative behaviors of real analysts. They typically use browser extensions or similar technology to observe how analysts handle incidents within their native interfaces (e.g., SIEM dashboards, case management tools).

The key value lies in their ability to transform institutional knowledge and best practices into scalable, reusable automation. Over time, these platforms can “replay” these expert workflows at scale, automatically handling new incidents just as a skilled analyst would, step by step, click by click.

This approach is particularly valuable for organizations seeking to preserve and multiply the

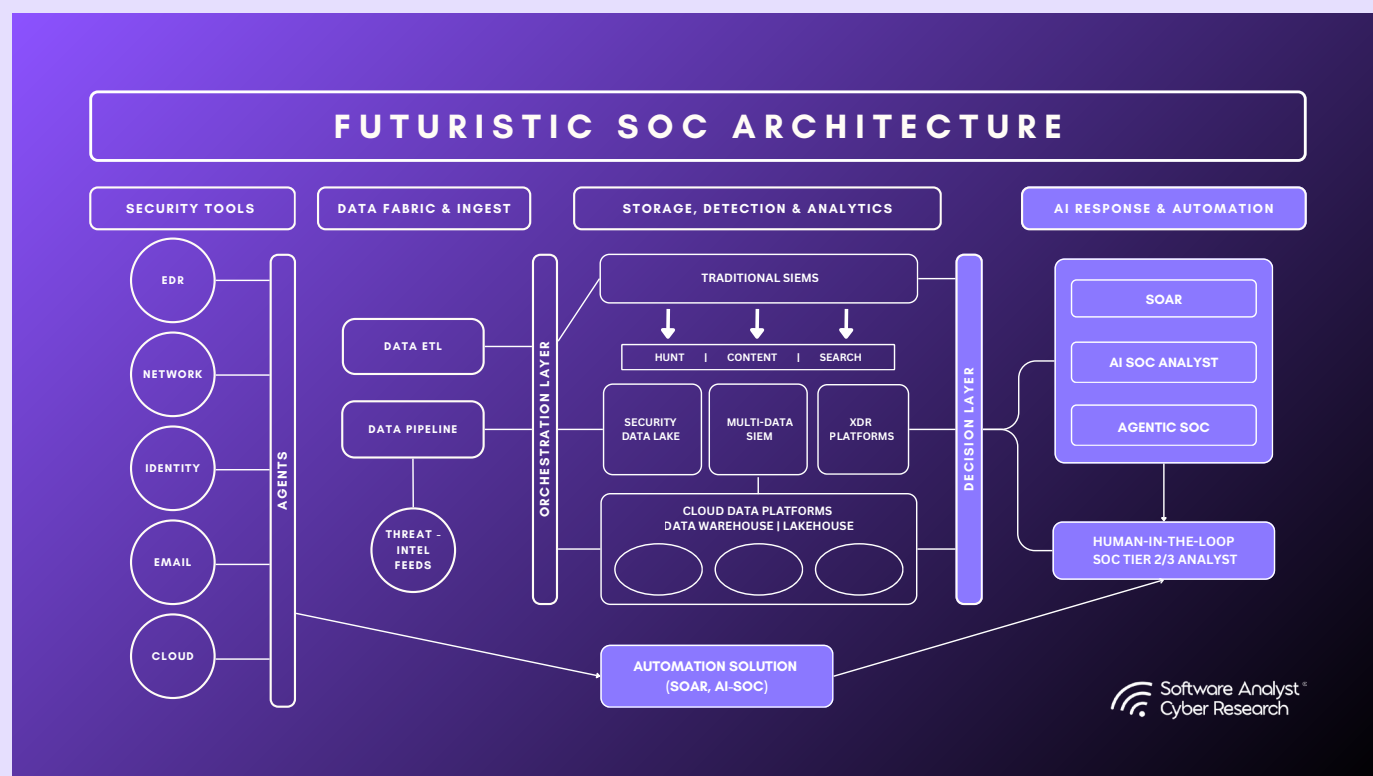
expertise of their best analysts, onboard new staff rapidly, or maintain strong consistency and quality across investigative processes.

However, there are some important caveats:

- These platforms require an upfront initial investment of time and expertise: workflows must be recorded and validated before value is realized.
- They may be slower to reach full operational impact compared to plug-and-play overlays.
- Their effectiveness depends on the presence of experienced analysts to “teach” the system

AI SOC vendors offer significant improvements in detection and response, but their solutions differ markedly in architectural design. Our evaluation identifies **three primary architecture types**, each with distinct strengths and limitations. Each architecture reflects a different approach to augmenting SOC operations. Some focus on enhancing what is already in place, while others aim to rebuild the investigative stack from the ground up. The right choice depends on an organization's maturity, goals, and existing toolset, as well as how much control and visibility they want over their data and detection logic.

Operational Guide to AI SOC Adoption



The following steps are designed to help organizations successfully integrate AI SOC products into their existing security workflows. This guide builds on our broader research to clarify which feature sets matter, identify vendors that align with specific operational needs, and provide context on where the market is heading. Some phases, particularly defining the AI strategy and selecting the right vendor, require thoughtful internal alignment and should be guided by rigorous proof-of-concept (POC) evaluations.

AI adoption is not a single decision. It is a staged process that benefits from cross-functional input, clearly defined success metrics, and a measured approach to automation. The phases below outline a practical path from early planning through full deployment, highlighting the operational checkpoints that matter most along the way.

Define the AI Security Strategy

We recommend that organizations begin by establishing a clear strategy for how AI will enhance their security operations. This should include identifying key pain points, such as alert fatigue or extended response times. AI objectives must be aligned with overarching business and security goals. Organizations should define success metrics upfront and secure stakeholder buy-in. Key questions to guide this phase include: Which SOC functions should AI improve or automate? What level of automation is appropriate? And how will success be measured?

Agree on the AI Feature Set

Next, organizations should define the capabilities required from an AI SOC solution. We recommend prioritizing core features such as threat triage and enrichment, AI-assisted investigation and response, behavioral analytics and anomaly detection, threat intelligence integration, and natural language querying. It is essential to engage stakeholders early to ensure alignment with operational needs. This is the stage where technology research and solution evaluation are most valuable.

Select an AI SOC Vendor

With requirements defined, organizations should select a vendor that aligns with their strategic goals and desired feature set. Solutions should be evaluated based on accuracy, explainability, and integration capabilities. Additional factors to consider include vendor support, regulatory compliance, and ease of deployment. Where possible, organizations should conduct pilot programs or proof-of-concept (POC) testing to validate solution effectiveness in real-world scenarios.

Deploy the AI SOC Solution

Following vendor selection, the focus should shift to integrating the AI solution into the existing environment. This includes connecting it with systems such as SIEM, EDR, SOAR, and telemetry sources. It is important to configure user roles, define operational workflows, and set up automation triggers. When available, historical data should be used to tune AI models for improved performance. Collaboration between internal teams and vendor support is critical during this phase.

Establish a Trust Period (1–2 Months)

Once deployed, we recommend an initial trust-building period focused on validating the AI system's performance. During this time, security analysts should closely review AI-generated alerts and decisions. Feedback loops must be implemented to continuously improve accuracy. Monitoring for false positives and making configuration adjustments is essential. The goal is to build confidence in AI's ability to support and enhance SOC operations reliably.

Transition to Full Automation

As trust in the system grows, organizations can gradually expand the level of automation. This may include automating alert triage and responding to low-risk incidents. With operational tasks increasingly handled by AI, security analysts can focus on strategic initiatives such as threat hunting, red teaming, and strengthening architecture and controls.

AI SOC can significantly enhance your organization's ability to detect and respond to threats. However, successful adoption requires a thoughtful approach to avoid choosing the wrong solution or overestimating what AI can deliver. We recommend a careful evaluation process, particularly before enabling automated actions. Our research shows that fully autonomous operation is achievable but only after thorough due diligence and a period of trust-building between your team and the AI SOC vendor.

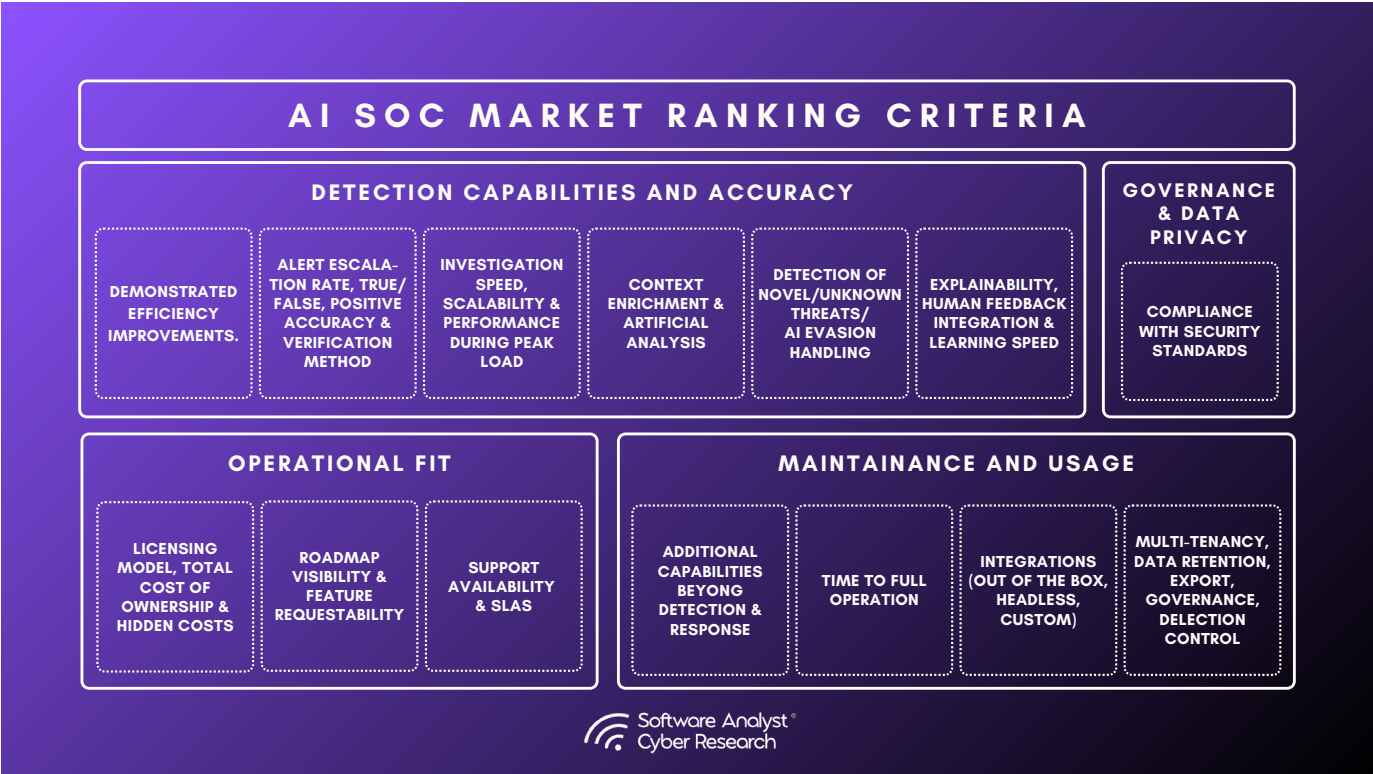
Risks and Considerations with Agentic AI in the SOC

Before diving into the methodology of this research, it is important to highlight the inherent risks of relying on Agentic AI solutions for SOC use-cases. The following considerations focus on business, operational, and compliance impacts, not technical limitations like model performance, integration effort, or feature sets.

These are the areas that leaders and decision-makers should review carefully before moving forward with agentic AI in the SOC use-cases.

- 1. Lack of Standardized Benchmarks:** There are currently no widely accepted benchmarks to evaluate agentic AI solutions in SOC environments. This makes it hard to assess performance, accuracy, or return on investment across different vendors.
- 2. Disruption by Established Platforms:** Major security vendors may integrate their own AI agents directly into existing tools, potentially reducing the need for standalone agentic AI products.
- 3. Hype Versus Proven Impact:** Agentic AI is a hot topic, but not all offerings deliver meaningful or measurable value. A cautious, evidence-based approach is needed when evaluating solutions.
- 4. Limited Differentiation Among Vendors:** Many vendors advertise the same core capabilities: triage, response, and explainability, making it difficult to distinguish real innovation from marketing noise.
- 5. Accountability and Liability:** Organizations need clear policies around responsibility and escalation when an AI system makes a wrong or harmful decision.
- 6. Compliance with Data Regulations:** Vendors must ensure that data is stored and processed in line with regional laws, including requirements around data residency and sovereignty (e.g., GDPR).
- 7. Changing Role of the Analyst:** Agentic AI shifts the SOC operating model. Analysts are moving from direct responders to overseers of automated systems, a transition that may require new skills, training, and changes to team structure.

Evaluation Methodology



To evaluate the capabilities of AI SOC vendors, we developed a structured assessment methodology based on real-world operational needs and expert input. Each vendor was measured across a comprehensive set of criteria that reflect the most critical functional, technical, and operational aspects of AI-driven security operations.

Our team collected data through vendor interviews, detailed questionnaires and product documentation. This information was then mapped against our evaluation matrix.

Below, we outline the key elements (not all) that make up our assessment matrix, along with explanations for why each factor is critical in evaluating the effectiveness and maturity of AI SOC platforms. These criteria reflect the capabilities we believe are most important for delivering real value in modern security operations.

1. Demonstrated efficiency improvements

While improvements in metrics are important, the method of tracking those improvements is even more crucial. We prioritized numbers from production environments and case studies over vague, high-level statements. To demonstrate platform value, vendors should provide data points such as:

Improvement of SOC Metrics: We assessed how automation reduced the time to acknowledge and investigate alerts. Could a full investigation be completed faster with the platform than with a human analyst? By how much? Vendors were expected to support claims with real-world numbers and case studies drawn from production environments.

Alert Escalation Rate: We reviewed how many alerts were fully handled by the platform versus those that still required human escalation. This ratio served as a proxy for the level of automation maturity.

Verdict Accuracy: We evaluated how each platform measured and maintained accuracy. In particular, we looked for structured QA practices rather than simple re-testing or analyst overrides.

The rigorous quality assurance (QA) on verdict accuracy was also essential. We gave lower weight to platforms that relied on AI re-checking its own outputs or client analysts confirming/denying verdicts. Human QA teams, detection engineers, and red team exercises where the AI SOC was evaluated, were key factors for scoring well in this category.

2. Investigation speed, scalability and performance under load

This assessment category is particularly relevant for large Managed Security Service Provider (MSSP) use cases, rather than internal, well-optimized environments. We assessed on the following:

Average Investigation Time: We took into account that investigation speed can be influenced by external factors, such as the responsiveness of APIs providing contextual information. Our focus was on determining the average runtime of investigations.

Scalability: We sought to understand the approach to scalability. Was it based on a fixed value of provisioned compute resources per license, or does the architecture offer inherent scalability? We requested vendors to explain how the product performs under unusually high peak load conditions.

Performance Under Load: We inquired whether the product had undergone stress testing and if there were case studies detailing such tests or large-scale deployments.

We valued transparent explanations regarding product architecture and how scalability was achieved. We recognize that while average investigation time is a consideration, ensuring support for peak loads and the queuing of alert investigations (rather than dropping them) in resource-constrained situations, are of greater importance..

3. Context enrichment and artifact analysts

We evaluated how well the platform gathers, links and analyzes additional information (such as file behavior, user activity, threat intelligence) to provide deeper context for investigations.

Contextual data sources: We looked for whether the platform came built in with threat intelligence components such as the ability to iterate alert entities against known databases of malicious artifacts – VirtusTotal, AbuseIPDB, Recorded Future etc.

Artifact analysis: We evaluated the depth of analysis on associated artifacts. This included use of built-in sandboxing, behavioral baselining, and UEBA-style techniques.

We valued vendors who brought integrations with popular services as part of their licenses. It was important for us to understand how deep the analysis goes. We valued sandboxing, UEBA-like approach to establishing baselines of what constitutes normal behavior and analysis methods with more depth than just comparing alert entities to lists of known bad artifacts.

4. Detection of Novel Threats

We are fully aware that AI SOC platforms are not detection engines on their own. Rather, these platforms focus on analyzing already generated alerts. Nevertheless, some platforms had features that allowed them to flag novel threats.

AI evasion: We examined how the platform defended against prompt injection and other forms of AI evasion across binaries, scripts, and log repositories.

Detection of novel threats: We assessed whether the platform included any capabilities to flag anomalous behaviors or unknown patterns not already surfaced by existing tools.

We valued vendors that demonstrated thoughtful approaches to AI evasion resistance and threat novelty.

5. Depth & Breadth of Integrations

We evaluated how well each platform integrated with the broader security stack and how smoothly it fit into analyst workflows.

List of available integrations: We reviewed whether major security related solutions could be integrated out of the box. Was SOAR included as well?

Headless mode support: We assessed whether the platform could operate without requiring analysts to log into a new console, what we refer to as “headless mode.” This reflects a realistic need, as many analysts already work across multiple platforms, and adding another interface introduces friction.

Integration with communication applications: We looked for the ability to query the platform using natural language from tools like Slack or Microsoft Teams, or through built-in chatbot features.

We valued integrations into major SIEM and case management tools. Tool fatigue is a valid concern, and not having to work directly within the AI SOC platform is a welcome feature. Lack of integration with major security vendors or SOAR workflows was considered a red flag. Given how straightforward it is to connect with modern SaaS-based security solutions, we expected these integrations to be fully implemented. Integration with communication applications was also a welcome addition developed by some vendors, not essential, but certainly a strong bonus.

6. Data & Privacy

We evaluated how each platform handled data control, governance, and multi-tenant environments, all key concerns for both MSSPs and large enterprises.

Support for multi-tenant deployments: We assessed whether the platform could effectively support multiple tenants, such as MSSPs managing distinct client environments or large enterprises with segmented business units.

Control over data: We reviewed the level of control customers had over their data, including retention settings, exportability, data deletion during offboarding, and the ability to choose storage locations. We also assessed how clearly vendors articulated their data-at-rest storage practices.

Compliance with security standards. We verified alignment with regulatory and industry frameworks such as ISO 27001, SOC 2, GDPR, and NIST. Auditability and documentation were also considered.

We placed high value on flexibility and transparency in managing data. Platforms that allowed data storage in the customer's own environment received added credit.

7. Explainability

We examined how clearly the AI system explained its actions and how effectively it incorporated analyst feedback into future decisions.

Explainability: We assessed whether every step in the AI-driven investigation was visible and auditable. This included checks for hallucination detection and clarity of decision paths.

Human feedback integration: We reviewed whether analysts could modify verdicts or influence the investigation logic, not just the final outcome. We also asked how quickly that feedback loop was implemented and how it affected future investigations.

All vendors provided some degree of step-by-step transparency. However, we prioritized platforms that enabled granular analyst input and rapid incorporation of that feedback. The more influence analysts could exert over the investigation process, the stronger the score.

Additional features

This category evaluated broader platform differentiation beyond core triage and investigation improvements. While most vendors demonstrated similar gains in Mean Time To Respond (MTTR) and investigation quality, we paid close attention to standout features and unique value propositions.

We assessed capabilities that extended beyond standard SOC workflows, such as support for compliance workflows, innovative analyst interfaces, or domain-specific use cases. These often revealed how well the platform could serve specialized teams or adapt to future needs.

We also evaluated practical operational factors like time to full deployment, vendor roadmap visibility, and customer support models. While these did not factor as heavily into the core capability scoring, they played a meaningful role in shaping the conclusions of this report.



R**DIANT**

Radiant Security

Radiant Security, founded in 2021 by former Exabeam executives, is a cybersecurity company focused on automating Security Operations Center (SOC) triage and response through artificial intelligence. Its core offering is an AI-driven SOC platform that seeks to address alert fatigue, reduce manual investigation workloads, and improve response time to security incidents. Radiant's system is designed to triage any security alert (including SIEM alerts) by integrating with existing enterprise environments and ingest telemetry from multiple sources, including alerts, endpoint logs, identity platforms, and any unstructured security telemetry and documentation.

The platform operates via a fleet of AI agents that perform automated triage and investigation across a wide array of security use cases. Its principal mechanism involves contextual enrichment of alerts, where extracted data artifacts are dynamically correlated with environmental signals such as user roles, behavioral baselines, and threat intelligence. Radiant employs just-in-time behavioral baselining rather than precomputed ML models, enabling investigations to adapt to context-specific activity patterns without continuous background processing. These baselines are generated from recent data, typically ranging from two weeks to 90 days, depending on the nature of the event.

Radiant distinguishes between known and unknown threats using a tiered model. For novel or previously unseen alerts, Radiant conducts ad hoc analysis and constructs new triage procedures, effectively augmenting the system's capabilities over time. Once analyzed, alerts are categorized into one of four statuses: benign, recommended benign, recommended malicious, or malicious. Only alerts deemed malicious or confirmed by a user are escalated to case-level incidents, helping reduce operational noise.

A significant focus is placed on automation. Radiant supports full-cycle automated remediation workflows, which can be configured to include or exclude human approval steps. This hybrid model allows organizations to gradually shift toward increased automation as trust in the system grows. Radiant's internal benchmarks suggest that full triage can occur within minutes, though mean time to respond (MTR) is often limited by user validation time. Customers reportedly reduce response cycles from industry-standard durations of several days to under eight hours, with automation-capable environments achieving even faster turnaround of as short as a few minutes.

Radiant's deployment model supports customer-controlled log storage, typically through Amazon S3 buckets as well as query connectors that leverage existing SIEMs and APIs. This configuration allows for compliance-friendly, long-term retention without the cost burden of traditional SIEM licensing models. The platform includes a built-in log management interface with support for Lucene queries and Grafana integration. This approach enables visualization and correlation without proprietary query languages or infrastructure dependencies.

Licensing is decoupled from data volume and is instead based on the number of users or alert types under management. This model aims to provide predictable pricing while scaling with organizational complexity and coverage requirements. For managed security service providers (MSSPs), Radiant offers licensing aggregation across tenants, with multi-bucket support to preserve data isolation and ownership.

From a customization standpoint, Radiant offers three primary levers: static allow/deny lists, environment-specific unsupervised learning, and codified natural language policies. These options

allow organizations to tailor alert disposition logic to their operational context and reduce false positives, particularly in environments where technical teams may generate anomalous but legitimate activity.

Radiant positions itself as a product-centric platform rather than a managed service. Customers are responsible for deployment and operation, with support limited to technical guidance. Initial implementation typically includes a four to six-week proof of concept, followed by a three-month “time to trust” period during which automation is incrementally enabled.

In summary, Radiant Security delivers an AI-based SOC automation platform focused on reducing alert volume, expediting investigation, and supporting scalable automation. Its emphasis on adaptive triage, just-in-time baselining, and customer-owned data infrastructure offers a distinct model within the AI SOC landscape.

Conclusion

The contemporary security operations center (SOC) has reached a critical juncture, driven by increasing alert volumes, persistent staff shortages, and growing cyber threat complexity. AI-driven solutions have become essential for alleviating these operational pressures by automating alert triage, improving investigative efficiency, and enabling scalability without significant headcount increases. However, successful adoption hinges on selecting architectural models—such as Overlay, Integrated Platforms, or Human Workflow Emulation—that align closely with organizational needs, maturity, and resource availability. Additionally, the transparency and explainability of AI-driven decisions have emerged as critical factors for gaining trust and ensuring effective integration into existing workflows. While the market for AI in SOC operations is expanding rapidly, it remains relatively immature. Organizations must therefore adopt AI solutions thoughtfully, balancing anticipated benefits against potential risks including compliance, data governance, and operational disruptions.

AI technologies are increasingly vital to modern SOC operations, offering practical benefits in automation, scalability, and response speed. Despite these advantages, organizations need to carefully evaluate the suitability of different AI architectures, demand transparency from solution providers, and manage implementation risks proactively. As the AI SOC ecosystem continues to develop, informed and cautious adoption will remain essential for effectively addressing contemporary cybersecurity challenges.

SACR will continue tracking this fast-moving space, with version 3 of this report set to come out soon. Stay tuned for deeper insights, expanded vendor coverage, and updated guidance as the market evolves.



business

personal



Trusted research. Sharp insights. Real conversation.

CISO

VENDOR

SECURITY
TEAMS

INVESTORS